

REMARKS/ARGUMENTS

Claims 1-28 were pending when last examined and are rejected. No claims have been amended and new claims 29-34 have been added. Therefore, upon entry of this amendment, which is respectfully requested, claims 1-34 will be pending. No new matter has been added.

Support for the amendments to the claims may be found in the specification as originally filed as follows. Support for the amendments to claims 1 and 11 may be found at least in the Abstract and paras. 0013, 0040-0043 and 0064. Support for new claims 29-34 may be found at least in paras.: (29) 0012, 0043, 0045, 0049 and 0056-0065; (30-32) 0033, 0039, 0043, 0061 and 0056-0065; (33-34) Abstract, 0011-0013 and 0040-0043.

Claim Rejections - 35 USC § 103

Brown i.v.o. Guthrie

In items 1-2 (page 2) of the present Action, the Examiner rejects claims 1-25 under 35 U.S.C. 103(a) as being unpatentable over Pub. No. US2003/0088579 to Brown et al. (hereinafter, "Brown"), and further in view of U.S. Patent No. 6,587,854 B1 issued to Guthrie et al. (hereinafter, "Guthrie"). However, the Examiner also specifically rejects claims 26-28 over Brown in view of Guthrie. It is therefore presumed by Applicants that the Examiner erred in stating the initial rejection in items 1-2, which instead assert a rejection of claims 1-28. Applicants respectfully traverse.

Claim 1:

The Examiner, on page 2, rejects claim 1, asserting that Brown teaches: a method of optimizing a query in a database comprising "generating tenant-level statistics for one or more of said plurality of tenants for one or more of the data table" and "generating tenant-level statistics for one or more of said plurality of tenants for one or more of the data table" at para.

0043; "receiving a SQL query" at para. 0033; and "optimizing the SQL query based on the statistic" at para. 0033-0034.

Applicants respectfully disagree with the Examiner's assertions regarding at least the Brown references and regarding establishing *prima facie* obviousness in view of Guthrie. For example, Applicants respectfully submit and the Examiner later at least impliedly admits that Brown 0043 does NOT in fact teach "generating tenant-level statistics..." for tenants in one or more tables of a multi-tenant database or otherwise. Contrary to the Examiner's initial assertion, Brown 0043 merely discloses a "*faster mechanism*... for collecting statistics of columns of a table in a database system." Specifically, rather "than collect[ing] statistics based on a full table scan..., statistics are collected based on reading a sample... of a table," which Applicants admit provide *faster* collection of table statistics than the full table scan that is replaced. Moreover, Brown 0033-0034 does NOT teach optimizing the SQL query based on a tenant-level statistic, but rather that a cost model uses [the para. 0043 table-level sample] statistics... to reduce the search space and thereby identify a query plan with a lowest response time (0033), and that the query plan is stored in a table of a query capture database (0034).

As noted earlier, the Examiner also later *admits* that Brown "does not explicitly [teach that] the database is a multi-tenant database." The Examiner therefore also at least implicitly admits that the Brown statistics of paras. 0044 and 0033-0034 can NOT be *tenant-level* statistics. Specifically, by admitting that Brown fails to teach a multi-tenant database, the Examiner necessarily also admits that Brown fails to teach multi-tenant database *tenants*, further, *tenant-level* statistics, and still further, optimizing according to *tenant-level* statistics as well – contrary to the Examiner's initial assertion Applicants further submit that no aspect of a multi-tenant database, tenants or tenant-level statistics are even considered, let alone rendered obvious in the cited references or elsewhere in Brown. The Examiner, however further asserts that Guthrie col. 4, line 66 – col. 5, line 15 (4/66-5/15) teaches "a multi-tenant database having one or more data tables, each table having one or more logical columns defining data categories and one or more logical rows associated with one or more tenants, wherein a plurality of tenants have data stored in the data tables."

Applicants presume arguendo that the Examiner asserts the following: Brown teaches table-level statistics in conjunction with a common database table and Guthrie teaches a multi-tenant database. Therefore, the Examiner reasons, it would have been obvious to one skilled in the art to replace Brown's random samples of table-level statistics with random samples of tenant-level statistics when operating in conjunction with Guthrie's multi-tenant database. (It "would have been obvious... to combine Brown and Guthrie... to optimize [a] query based on enterprise [i.e., tenant] level... statistic... in a multi-tenant database.") (Emphasis is added.)

Applicants respectfully disagree with the Examiner's assertions for at least the foregoing reasons, and at least because no incentive exists in Brown for the proposed modification, the present invention first recognized the problem of inaccuracies due to non-tenant-aware statistics collection (e.g., see instant Background) and the advantages of tenant aware statistics collection and optimization, and application of the present invention has since proven extremely successful. However, the point is rendered moot by way of the present amendment to claim 1 (which is entered in order to expedite prosecution).

Applicants nevertheless further submit that the Examiner's asserted basis for the combination is also erroneous and suggests impermissible hindsight. Specifically, the Examiner asserts that one... would "combine Brown and Guthrie ... in order to provide a more *accurate* method and *reduce the processing time* since only a subset of data associated with a particular tenant are processed instead of a whole table". Applicants submit, however, that Brown discloses using *random samples* of "database-level information" (0028), such as "the first M rows" or "every Nth row" (0050) as a "faster mechanism... for collecting statistics of columns of a table" "rather than... a full table scan" (0043-0044). Guthrie is silent in this regard, as the Examiner at least impliedly admits. Therefore, both Brown's reasoning and random samples are clearly NOT directed at and clearly do NOT teach or suggest and provide NO incentive whatsoever for a more "accurate" method, but are instead simply directed at the disclosed faster collection via random sampling of an ordinary database table. (Emphasis is added.)

Independent claim 1, as amended, recites:

“1. (Currently Amended) A method of optimizing a query in a multi-tenant database, said multi-tenant database having one or more data tables, each table having one or more logical columns defining data categories and one or more logical rows associated with one or more tenants, wherein a plurality of tenants have data stored in the data tables, the method comprising:

determining database indices for one or more of the data tables;
generating tenant-level statistics for one or more of said plurality of tenants for one or more of the data tables;
receiving a SQL query;
optimizing the SQL query based on the database indices; and
optimizing the SQL query based on the tenant-level statistics,
thereby enabling SQL query optimization according to greater semantic knowledge of use of the data tables.”

Applicants further submit, in addition to the above, that Brown in view of Guthrie fails to render obvious at least the recited “determining database indices for one or more of the data tables”; and “optimizing the SQL query based on the database indices.”

Applicants therefore respectfully submit that unamended claim 1 and claim 1 as amended are patentable over Brown in view of Guthrie for at the foregoing reasons and withdrawal of the rejection is respectfully requested.

Claim 2

On page 3 of the present Action, the Examiner further rejects dependent claim 2, which depends from claim 1. The Examiner asserts, in addition to “Brown and Guthrie teach the method of claim 1,” that the remainder of claim 2 is taught as follows: “wherein tenant includes one or more associated users” (Guthrie at 5/64-67); “generating user-level statistics for one or more of the users of one or more of the tenants for one or more of the data tables” (Brown

0043); and “optimizing the SQL query based on the user-level statistic” (Brown 0033).

Applicants respectfully disagree.

Applicants respectfully submit, in addition to the arguments already submitted respecting claim 1, at least that Brown fails to teach *or otherwise render obvious* at least “generating user-level statistics,” let alone “for one or more of the users of one or more of the tenants” or “for one or more of the data tables” at para. 0043 or elsewhere. Brown also fails to teach *or otherwise render obvious* at least “optimizing the SQL query based on the user-level statistic” (at para. 0033 or elsewhere), the combination of tenant-level statistics of unamended claim 1 and user level statistics of claim 2, or the combination of database indices and tenant-level statistics of amended claim 1 with the user-level statistics of claim 2. (No references were cited respecting the combinations.) Moreover, the Examiner does not assert that Guthrie cures any of these defects in Brown and Applicants submit that it does not. Therefore, Brown in view of Guthrie also fails to render claim 2 obvious under 35 USC 103.

Brown para. 0043, for example, would tend to contradict the Examiner’s assertion and further teaches away from the invention and the Examiner’s assertion. Specifically, Brown 0043 discloses “a *faster* mechanism is provided for collecting statistics columns of a table in a database system. Rather than collect statistics based on a full table scan..., statistics are collected based on reading a sample... of a table. The sample is set by specifying a percentage... indicating the percentage of rows to read from the... table in collecting the statistics.” (emphasis added) Brown para. 0044 further discloses that “the collection of statistics is to be based on a sample of a... table (rather than the entire table).” (emphasis added) Clearly, Brown para. 0043 does NOT teach (or otherwise render obvious) the recited “generating user-level statistics...” or generating any combination of statistics, but instead, merely that collecting random sample of table-level statistics is faster than collecting statistics on the entire table (which the Examiner *already* asserts one skilled in the art would modify as tenant-level statistics “*instead of a whole table*”). Brown para. 0033 merely discusses identifying “a query plan that has the lowest, or one of the lowest, response times, a cost optimization model and that the optimizer uses statistics and sampling to reduce the search space, and clearly provides NO support whatsoever for the Examiner’s assertion or otherwise for supporting *prima facie* obviousness. (Emphasis is added.)

Moreover, Applicants respectfully submit that Brown teaches away from the recited embodiments and that modifying Brown to render obvious the recited “generating user-level statistics...” or to generate the any combination of statistics would impermissibly render Brown unsatisfactory for its intended purpose and would impermissibly change the principle of operation of Brown (MPEP 2143.02), among other deficiencies. Specifically, Brown repeatedly discloses that its purpose is to provide for “faster collection” and a “faster mechanism” for collecting table-level statistics (e.g., Abstract, cited paras. 0033, 0043, 0044) according to the principle that collecting random samples of table-level data (0028), such as “the first M rows” or “every Nth row” (0050) of columns of a table “rather than... a full table scan” (0043-0044) will provide such a “faster mechanism... for collecting statistics (0043-0044). Thus, the asserted additional collecting and optimizing of user-level data or further not-considered combinations of data would clearly contradict the “faster” operation of Brown (even ignoring that Brown mentions users but does not even consider collecting or optimizing on user-level statistics, let alone some combination of statistics “levels” that is instead clearly gleaned from the instant invention). Further, adding additional processing as the Examiner asserts would render Brown unsatisfactory for its intended purpose of faster collection of table-level statistics, and would modify Brown’s principle of operation, i.e., achieving faster collecting of table-level data by collecting only lesser random samples of the table-level data instead of collecting statistics on all of the table-level data.

Applicants further respectfully submit that the Examiner’s assertion respecting claim 1, while providing the Examiner’s basis for citing Brown at all, clearly contradicts the Examiner’s assertion regarding claim 2. As discussed regarding claim 1, the Examiner’s only asserted basis for the present 103 rejection is that “in order to... reduce the processing time... only a subset of data associated with a particular tenant are processed instead of [processing] a whole table.” Thus, assuming arguendo that the Examiner is correct, with which Applicants disagree, then the asserted basis for replacing table-level statistics sampling with tenant-level statistics sampling (i.e., “instead of”) contradicts collecting/optimizing according to table-level statistics or both of table-level and tenant-level statistics in claim 1 as amended, further contradicts collecting/optimizing according to both of tenant-level statistics in unamended claim

1 *and* also collecting user-level statistics in unamended claim 1 and claim 2, and still further contradicts (if not, even more so) collecting/optimizing according to both of table-level statistics and tenant-level statistics in amended claim 1, or also collecting user-level statistics in claim 2. As discussed in the previous paragraphs, the Examiner's assertion is not only completely unsupported, but also contradicts at least one or both of the Examiner's asserted bases of speed and replacement (i.e., speed and replacement of table-level statistics sample collection).

Accordingly, Applicants respectfully submit that claims 1 and 2 are patentable over the combination of Brown in view of Guthrie for at least the foregoing reasons and withdrawal of the rejection of claims 1 and 2 is respectfully requested.

Claims 2-10

As submitted above, the Examiner's rejections of dependent claims 2-10, which depend from claim 1, and claims 3-6 further depending from claim 2, are also respectfully traversed and Applicants submit that claims 2-10 may include further features that are independently patentable over Brown in view of Guthrie. However, in order to expedite prosecution, specific aspects of claims 3-10, and other dependent or otherwise related or "asserted as related" claims, the rejections for which are hereby traversed, will not be fully argued in a separate manner. Accordingly, Applicants respectfully submit that claims 2-10 are dependent claims depending from claim 1 or further claim 2 and are patentable over Brown in view of Guthrie for at least the same reasons that claims 1 and 2 are patentable over Brown in view of Guthrie.

Withdrawal of the rejection of claims 2-10 is therefore respectfully requested for at least the foregoing reasons.

Claim 11-20

Continuing on pages 5 of the present Action, the Examiner rejected independent claim 11 and dependent claims 12-20 for "the same reasons" as with claims 1-10.

Applicants respectfully submit that claim 11, while independently patentable, includes similar limitations as with amended claim 1 and claim 2, and is patentable over Brown

in view of Guthrie for at least the same reasons that amended claim 1, and claim 2 are patentable over Brown in view of Guthrie. Claims 12-20, which depend from claim 11 and are also independently patentable in manners that may be independently argued, also includes similar limitations as with claims 1-10. Therefore, in order to expedite prosecution, Applicants submit that claims 11-20 are patentable over Brown in view of Guthrie for at least the same reasons that claims 1-10 are patentable over Brown in view of Guthrie.

Withdrawal of the rejection of claims 11-20 is therefore respectfully requested for at least the foregoing reasons.

Claims 21-25

The Examiner also rejected independent claim 21 and dependent claims 21-25 for “the same reasons” as with claims 1-10.

Applicants respectfully submit that claim 21, while independently patentable, includes similar limitations as with unamended claim 1, amended claim 1 and claim 2, and is patentable over Brown in view of Guthrie for at least the same reasons that unamended and amend claim 1, and claim 2 are patentable over Brown in view of Guthrie. Claims 21-25, which depend from claim 21, are also independently patentable in manners that may be independently argued and also include similar limitations as with claims 1-10. Therefore, in order to expedite prosecution, Applicants submit that claims 21-25 are patentable over Brown in view of Guthrie for at least the same reasons that claims 1-10 are patentable over Brown in view of Guthrie.

Withdrawal of the rejection of claims 21-25 is therefore respectfully requested for at least the foregoing reasons.

Claims 26-28

Continuing on pages 5 of the present Action, the Examiner also rejected claims 26-28, relying –at least in part– on the rejection of claims 1 and 11.

Applicants respectfully submit that claims 26-28, which depend from respective ones of claims 1, 2 claim 11, are independently patentable in manners that may be independently argued. However, in order to expedite prosecution, Applicants submit that claims 26-28 are

patentable over Brown in view of Guthrie for at least the same reasons that claims 1, 2 and 11 are patentable over Brown in view of Guthrie.

Withdrawal of the rejection of claims 26-28 is therefore respectfully requested for at least the foregoing reasons.

Response to Arguments

In item 3 (page 6) of the present Action, the Examiner responded to Applicants' arguments in the previous, non-final office action. Applicants respectfully object to the Examiner's interpretation of the respective terms at least as contradictory to the very references that the Examiner asserts to have used in conducting interpretation of the respective terms, and as failing to properly interpret the claim terms as a whole or consistently with the specification under MPEP 2111. (While Applicants also respectfully disagree that Applicants failed to provide sufficient explanation of the terms, the requirements of MPEP 2111 do not appear to be averted for purposes of examination merely because the Examiner is dissatisfied with the sufficiency of Applicants' remarks.)

Applicants respectfully submit, for example, that the Examiner's asserted interpretation of "multi-tenant database" as "a database having one or more tables, accessible by two or more users" contradicts the Examiner's current *and later* assertions, and further fails to meet the requirements for claim interpretation as conducted by the Examiner. For example, the Examiner looks to the claims last, asserting that the claim 1 preamble recites "said multi-tenant database having one or more data tables, each table having one or more logical columns defining data categories and one or more logical rows associated with one or more tenants," which contradicts both of a mere "database" and such mere database that is "accessible by one or more users." Additionally, as the Examiner later admits in discussing "tenant-level statistics," the unamended claim 1 preamble instead recites "A method of optimizing a query in a multi-tenant database, said multi-tenant database having one or more data tables, each table having one or more logical columns defining data categories and one or more logical rows associated with one or more tenants, wherein a plurality of tenants have data stored in the data tables," which even more clearly contradicts the Examiner's former only-partial recitation and asserted "a database

having one or more tables, accessible by two or more users.”

Moreover, “multi-tenant database” embodiments are also taught in the specification at other than the “multi-tenant database system” descriptions of the initial “Background” paras. 0001 and 0002 that the Examiner instead elected to consider. The instant specification “Detailed Description,” for example, teaches that “According to one embodiment, each MTS 16 is configured to provide web pages, forms, data and media content to user systems 12 to support the access by user systems 12 as tenants of MTS 16. As such, MTS 16 provides security mechanisms to keep each tenant's data separate unless the data is shared” (para. 0028). Other relevant teachings are also provided at least at paras. 0029 and 0031-0037, 0041, 0043, 0045, 0050, 0051, 0053, and so on, as well as in the corresponding drawings. Each such reference would further contradict the Examiner’s assertion that a “multi-tenant database” is merely “a database having one or more tables, accessible by two or more users.”

On page 7, the Examiner further interprets “tenant-level statistics” as “statistics collected on a subset of rows in a table.” However, once again, the Examiner’s interpretation contradicts even the Examiner’s own asserted references/reasoning, let alone the teachings of the instant application or the requirements for claim interpretation. For example, the Examiner asserts reasoning that “in case the table having one or more rows associated with more than one tenants, the statistics are collected based on... only the rows associated with a particular tenant” and nevertheless concludes that “tenant level statistics” are properly interpreted as “statistics collected on a subset of rows in a table.” This also makes no sense. Moreover, this interpretation makes no sense in view, for example, of the claim 2 “user-level statistics” and the references in the specification at least at paras. 0012, 0013, 0015, 0016, 0045, 0047-0051, and so on.

Applicants therefore respectfully submit that no proper examination may be made on such clearly improper interpretation of the claims. (See also the Examiner’s application of such improperly asserted interpretation at pgs. 7-8 of the present Action.)

Applicants further respectfully object to the Examiner’s mis-characterization of Applicants’ previous Arguments as not only inaccurate, but also as being improper under at least

MPEP 2142, 2143 and 2111, even if Applicants' previous Arguments were presented as the Examiner asserts. For example, even assuming *arguendo* that Applicants, in arguing claim 1, merely stated that "Brown at paragraph 33-34 fails to even mention 'the statistics' (with which Applicants respectfully disagree), MPEP 2142, 2143 and 2111 would require the Examiner to consider claim 1 "as a whole." Thus, the Examiner is clearly required to consider the corresponding subsequent claim 1 element, i.e., "the statistics" (and thus the corresponding argument), to more properly refer to the claim 1 element forming the antecedent basis for "the statistics," i.e., "tenant-level statistics," which are indeed NOT even mentioned in Brown paras. 0033-0034.

Moreover, Applicants respectfully object to the Examiner's apparent assertion of yet another interpretation of Brown or prior failure to express such assertion in a clear manner that would enable Applicants to respond. Here, the Examiner –having mis-interpreted "tenant-level statistics" as merely "statistics collected on a subset of rows in a table (e.g., see above)- further apparently mis-interprets Brown as teaching the "sample" as meaning essentially the same thing, thereby again contradicting the Examiner's earlier assertion. (The Examiner now asserts: "Brown therefore teach most of the limitations of claim 1. Brown teaches the step of generating tenant-level statistic for one or more of said plurality of tenants and optimizing the SQL query based on the tenant-level statistics".)

Applicants submit that Brown does NOT teach sampling of specific rows of a table according to either the data characteristics of data stored in a the row or particular data characteristics that are not even considered by Brown in a database unlike those to which Brown is addressed, or further, additional information provided in a not-considered database type. Rather, as was discussed earlier, Brown teaches using random samples of "database-level information" (0028), such as "the first M rows" or "every Nth row" (0050) as a "faster mechanism... for collecting statistics of columns of a table" "rather than... a full table scan" (0043). It is these samples that Brown refers to in para. 0043 and the statistics derived therefrom that Brown refers to in para. 0033. With respect, the tendency of the Examiner to segment and dis-associate associated terms is again objected to by Applicants, and Applicants renew and hereby re-submit all arguments presented in the response submitted by Applicants in response to

the previous Office Action.

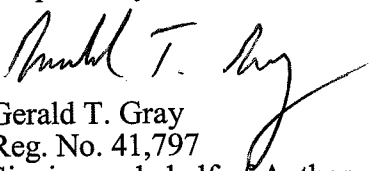
Applicants therefore further respectfully request withdrawal of the finality of the present Action for failure of the Examiner to meet the requirements for examination (e.g., failure to properly consider the claim terms, failure to properly interpret the claim terms), in further view of the substantial deficiencies exhibited by non-existent or substantially inaccurately asserted teachings in the cited references and the resulting rejection, and based on the Examiner's failure to properly consider Applicants' Arguments.

CONCLUSION

In view of the foregoing, Applicant believes all claims now pending in this Application are in condition for allowance. Withdrawal of the rejections, withdrawal of the finality of the rejections and the issuance of a formal Notice of Allowance at an early date is respectfully requested.

If the Examiner believes a telephone conference would expedite prosecution of this application, please telephone Daryl Josephson at (650) 631-4232.

Respectfully submitted,


Gerald T. Gray
Reg. No. 41,797
Signing on behalf of Author
Daryl Josephson, Reg. No. 37,365

TOWNSEND and TOWNSEND and CREW LLP
Two Embarcadero Center, Eighth Floor
San Francisco, California 94111-3834
Tel: 925-472-5000
Fax: 415-576-0300
GTG:sea
60817053 v1